

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026

Mastering Fabric Data Agents: From Setup to Success

Wednesday March 18th

FABCON | **SQLCON**

ATLANTA26

JOIN THE
CONVERSATION [#FABCONSQLCON26](#)

Personas

Customer

Financial Advisors

Portfolio Managers

...

Modern Technology

Co-pilot / Chatbot

Predictions

Data Management
Platform



Fabric

Platforms

Digital Experience

CRM & ERP

Document
Management

...



Marc Lelijveld

Technical Evangelist | Solution Architect
Macaw Netherlands



 [linkedin.com/in/MarcLelijveld](https://www.linkedin.com/in/MarcLelijveld)

 Data-Marc.com

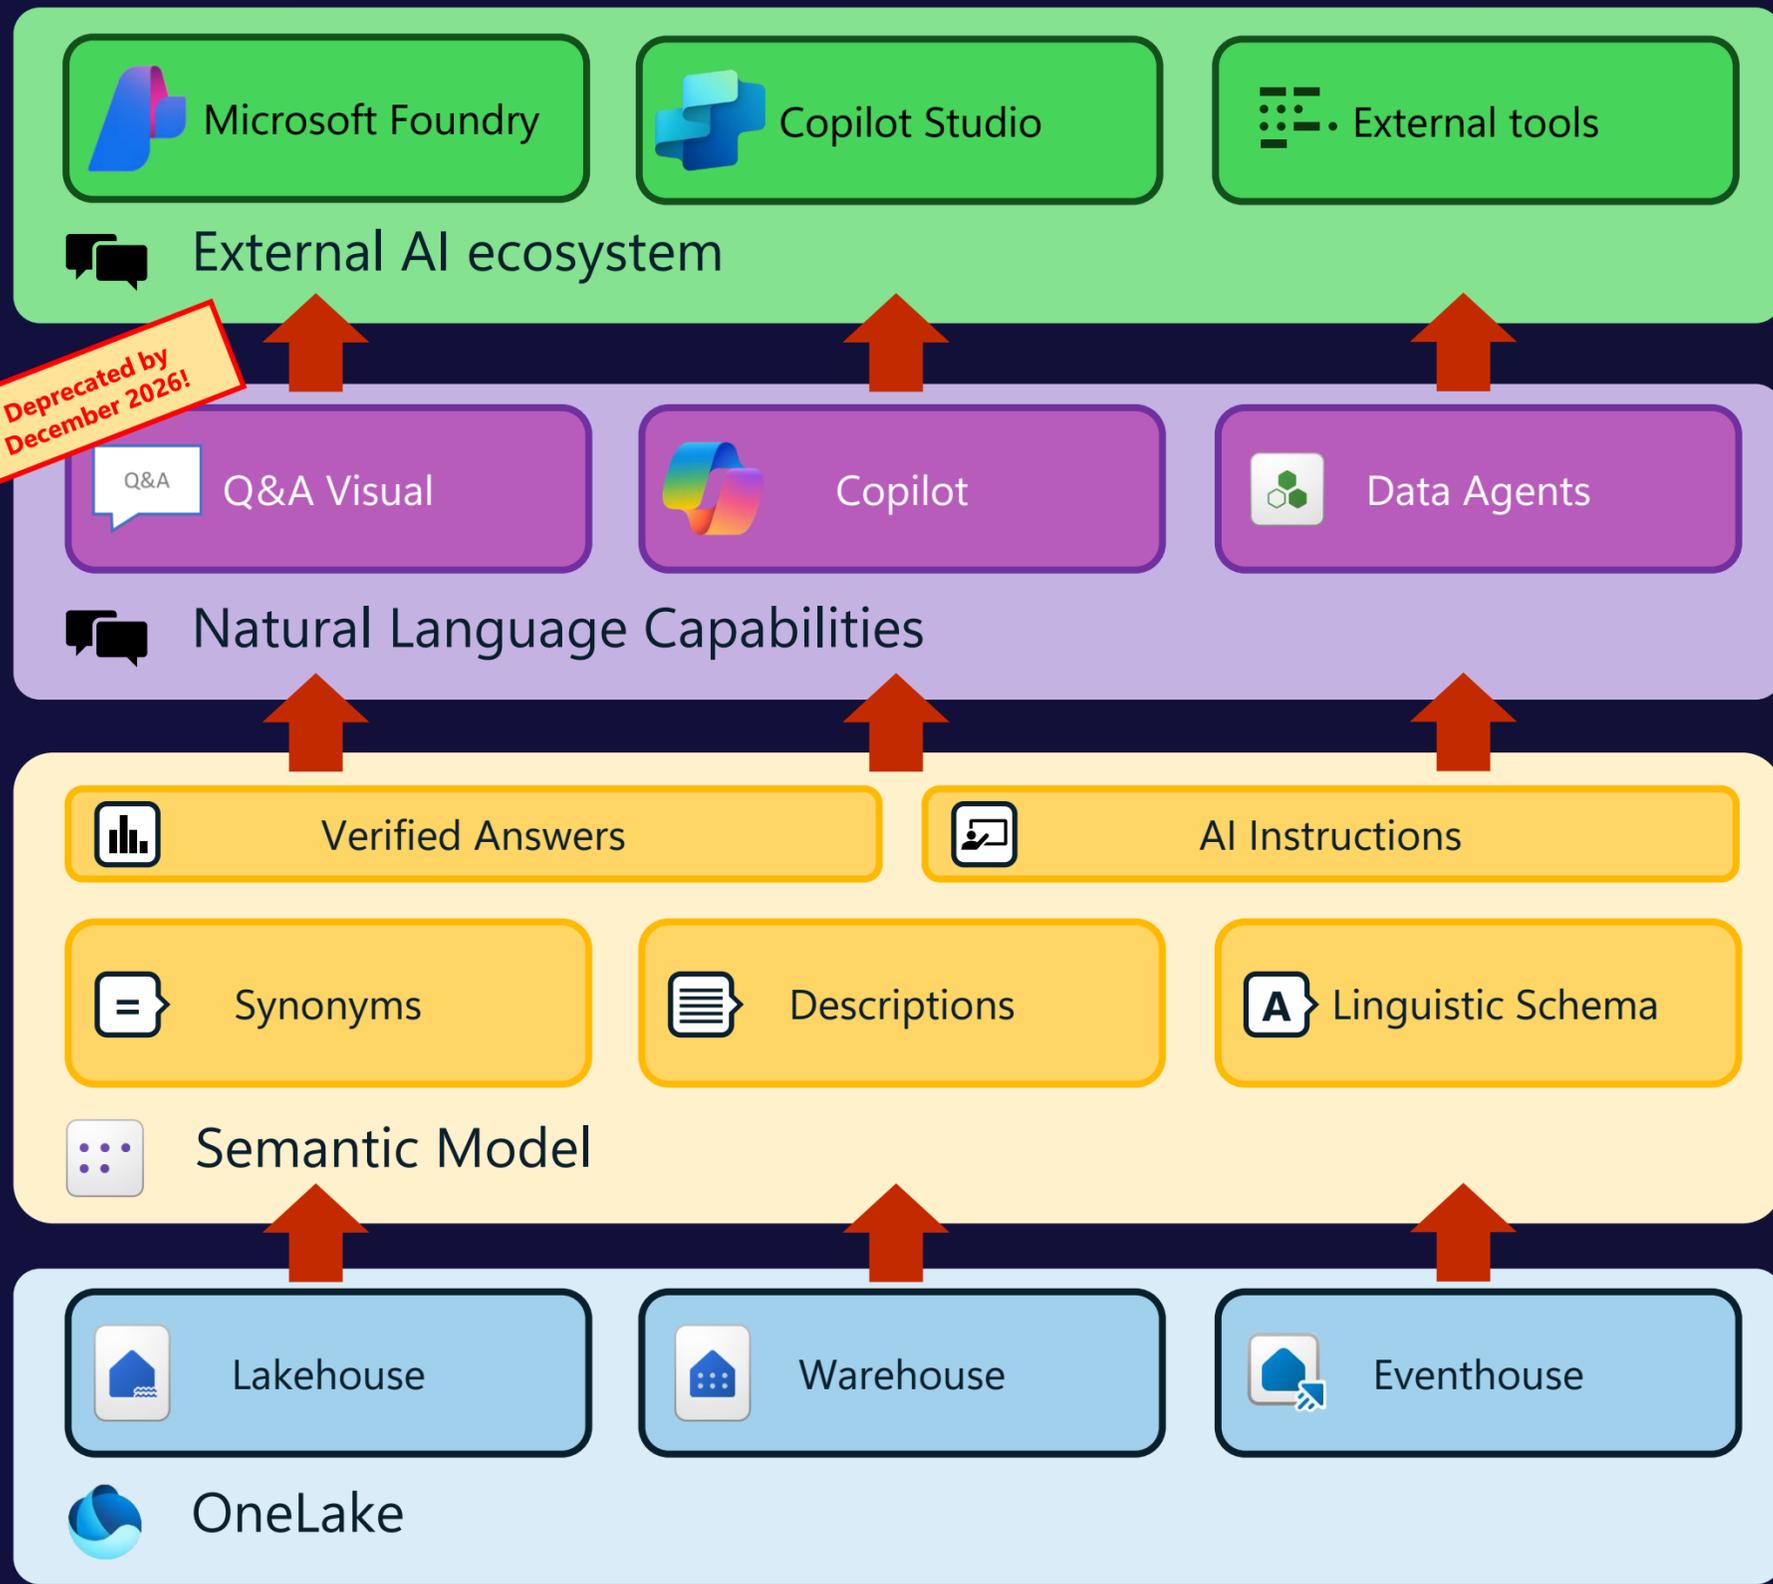


Exploring solution areas

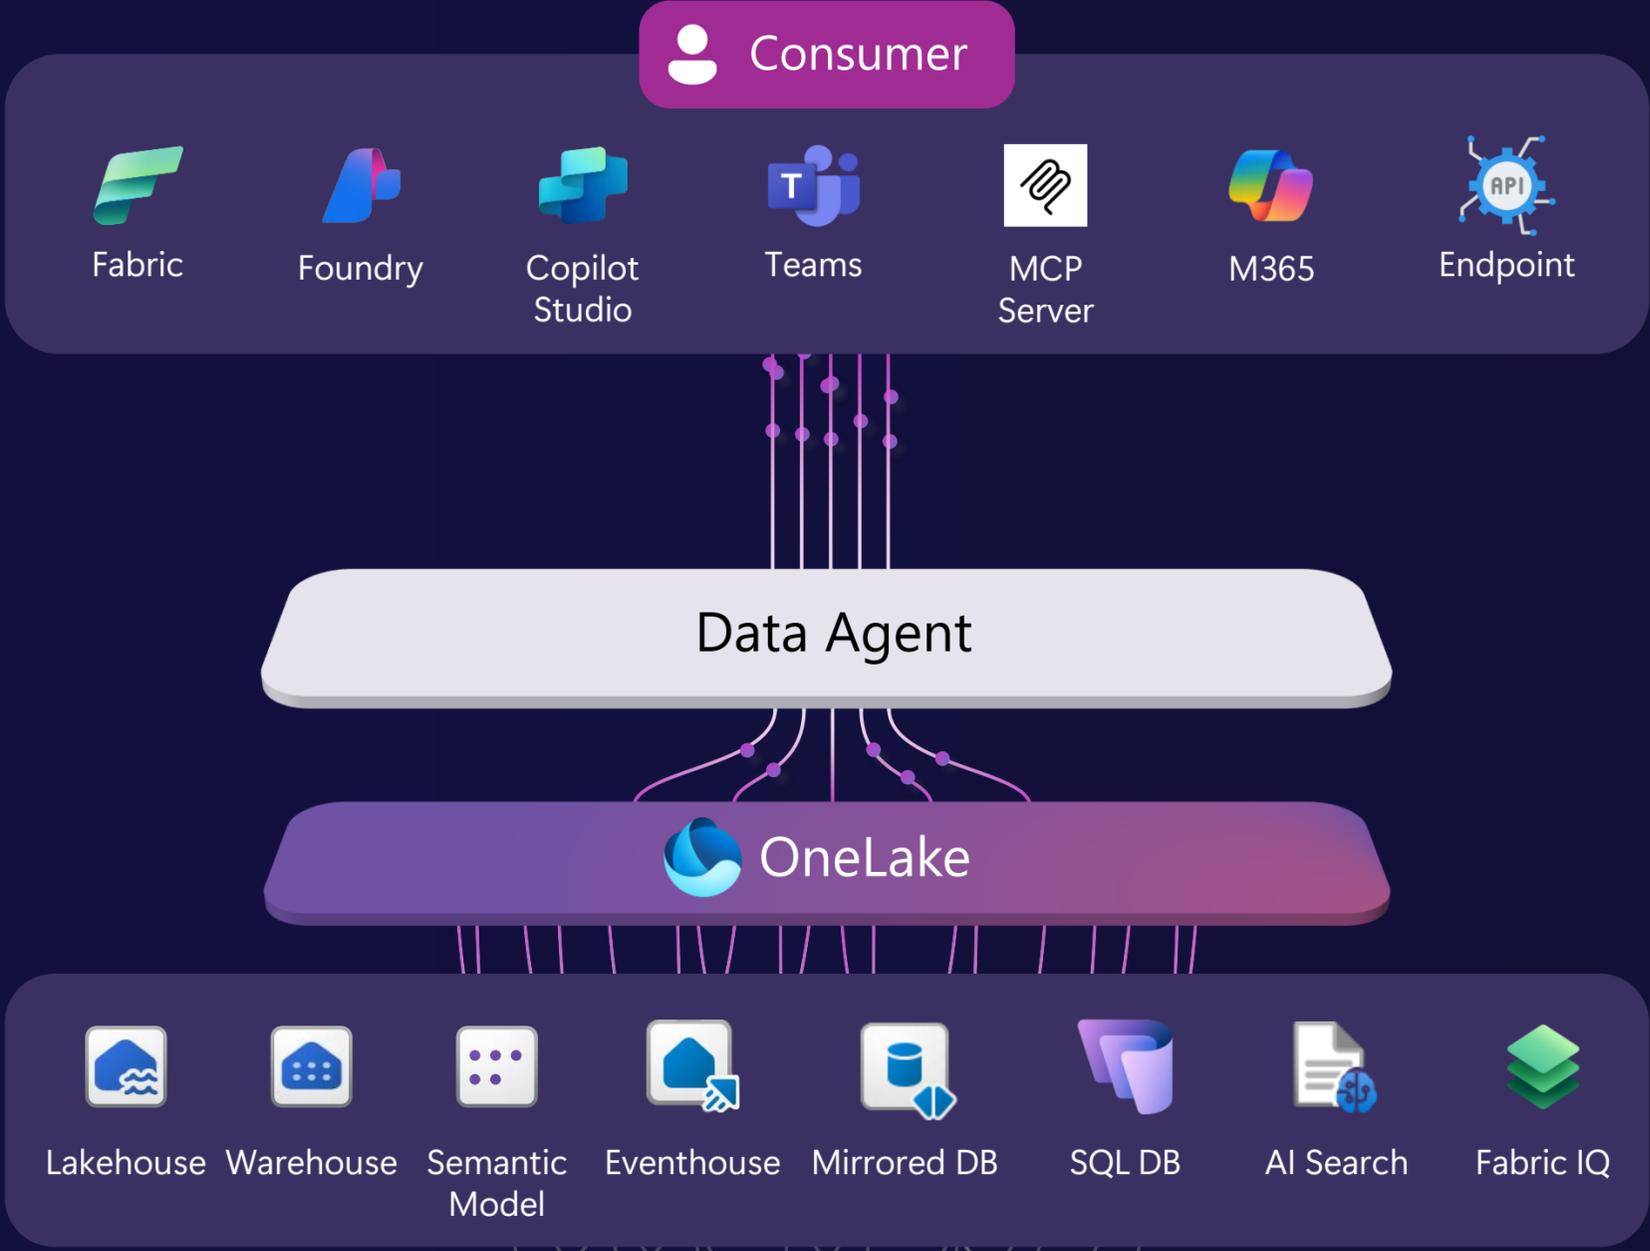
What are the alternatives to copilot?

What can you conclude from the orders placed versus the orders shipped?





Fabric Data Agents



A **virtual analyst** that allows users to interact with and gain insights from enterprise data in OneLake

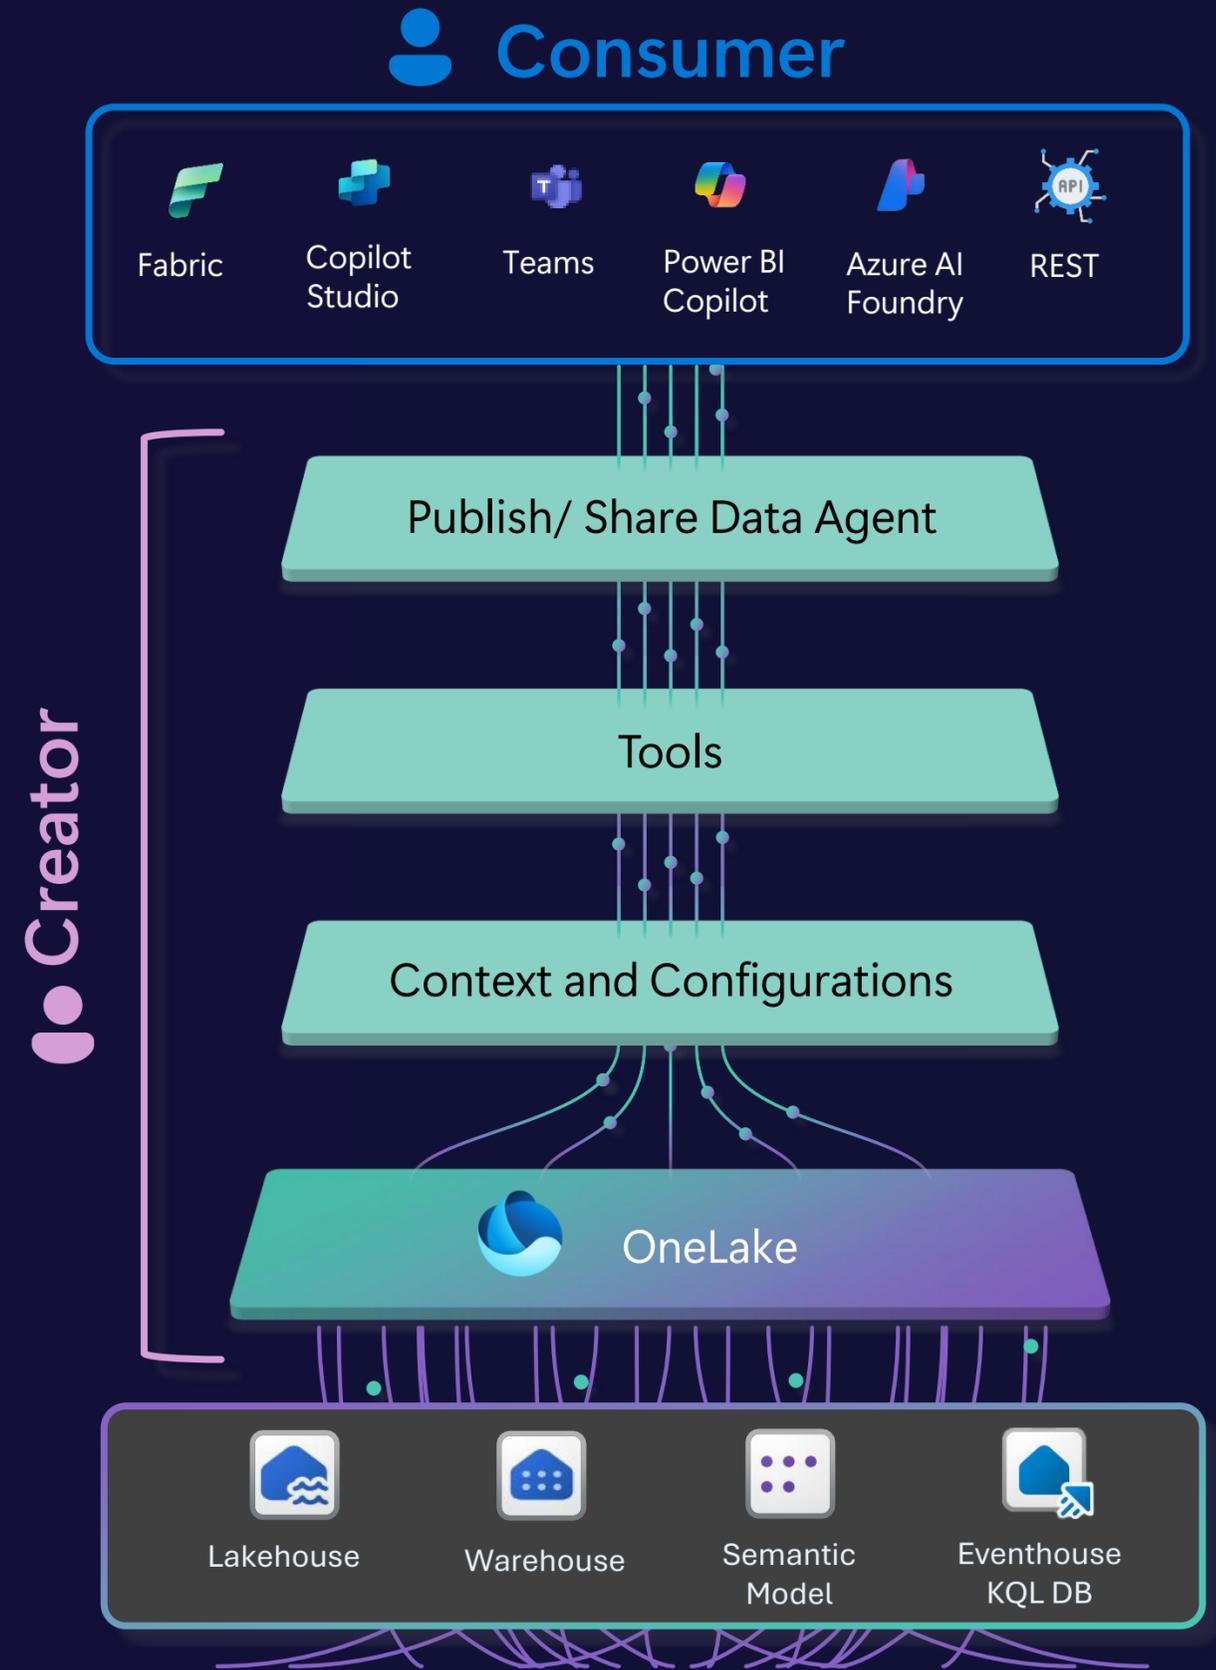
New Consume your data agents in **M365 Copilot** and as **remote MCP Server in VS Code**

New Data agents now support unstructured data through **Azure AI Search**

New Data agents now support **Fabric IQ Workload through Ontology**

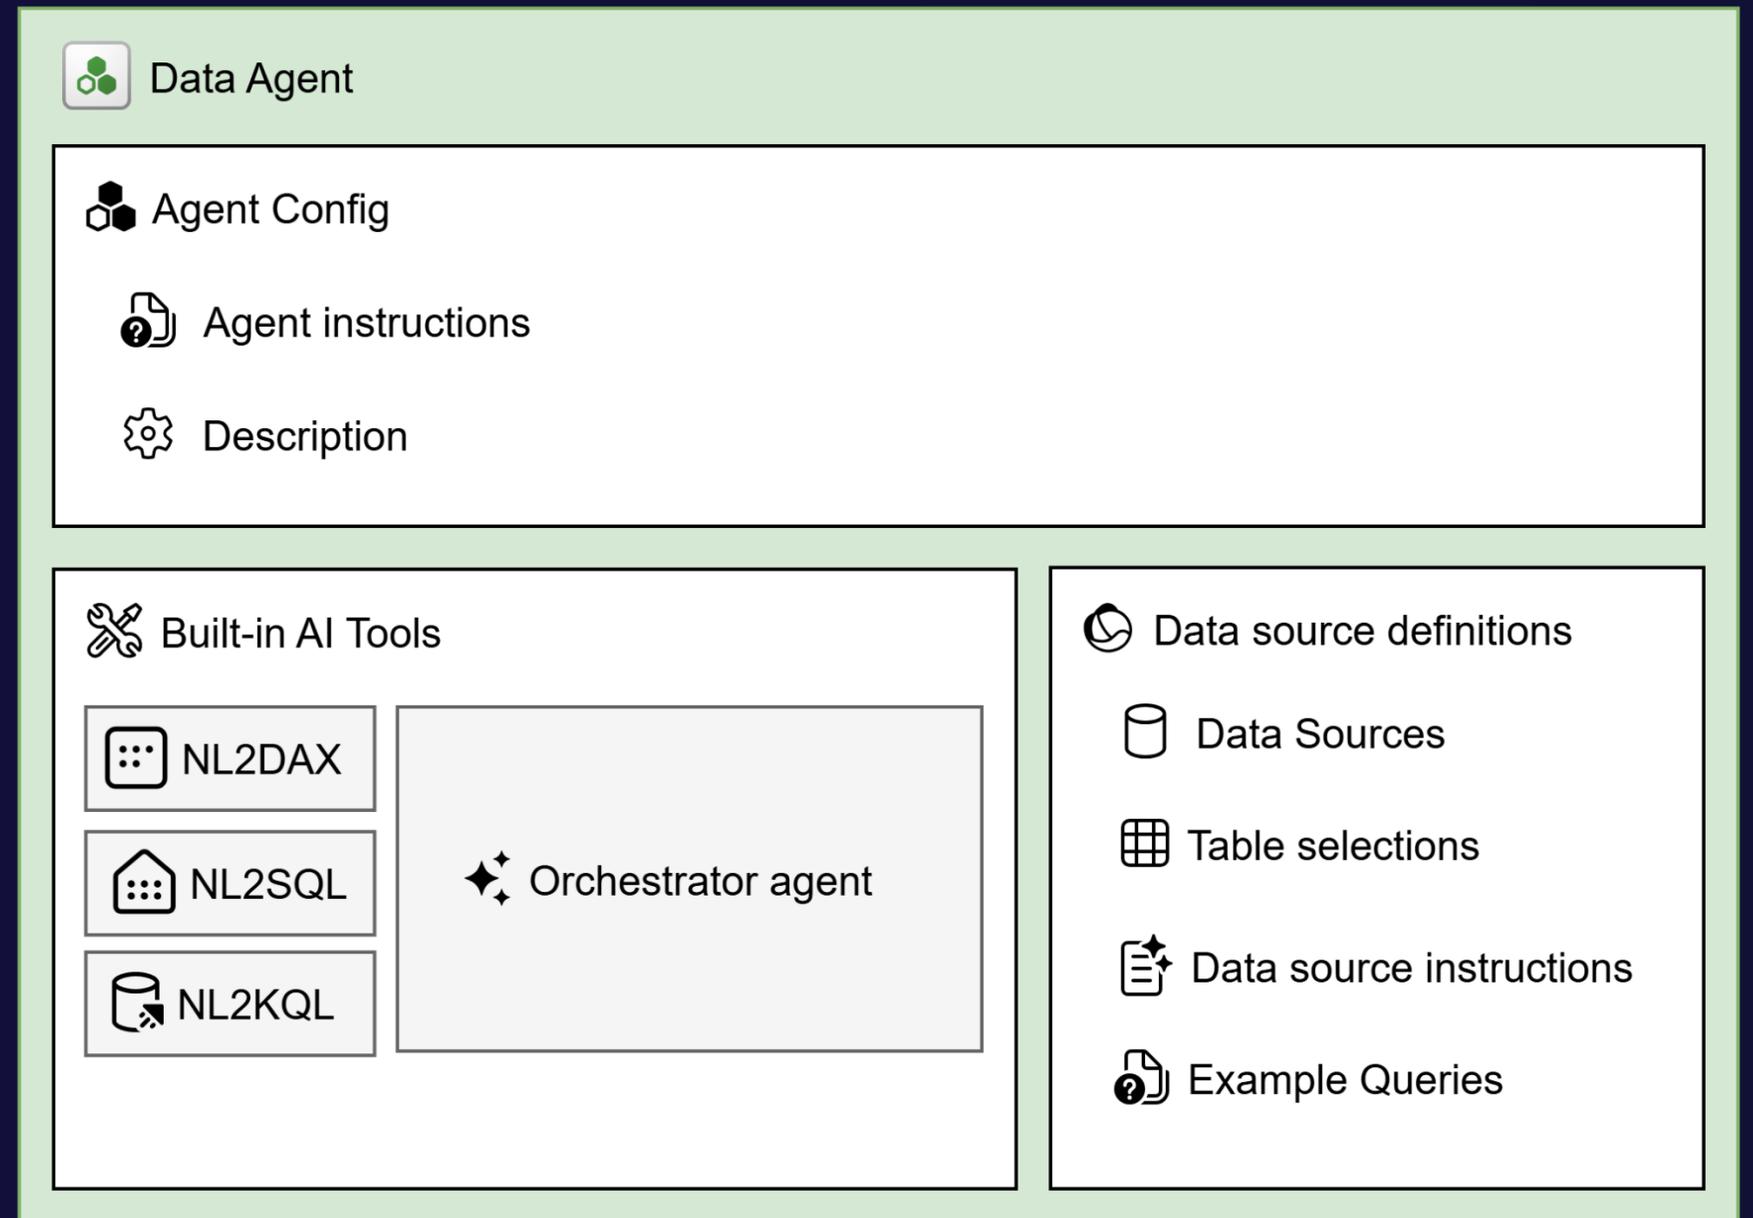
Introducing Fabric Data Agents

-  The **conversational Data Agent** allows users to interact with data naturally, enhancing accessibility and usability.
-  Improved chat canvas for creators with **new debugging capabilities**, making it easier to understand and refine responses.
-  Seamlessly reason over multiple data sources, including **Lakehouses, warehouses, semantic models and Eventhouse KQL databases**.
-  Your Data Agent can be consumed inside and outside of Fabric. Stay tuned for **upcoming integrations** with Copilot Studio, Teams, Azure AI Foundry and your own custom applications.

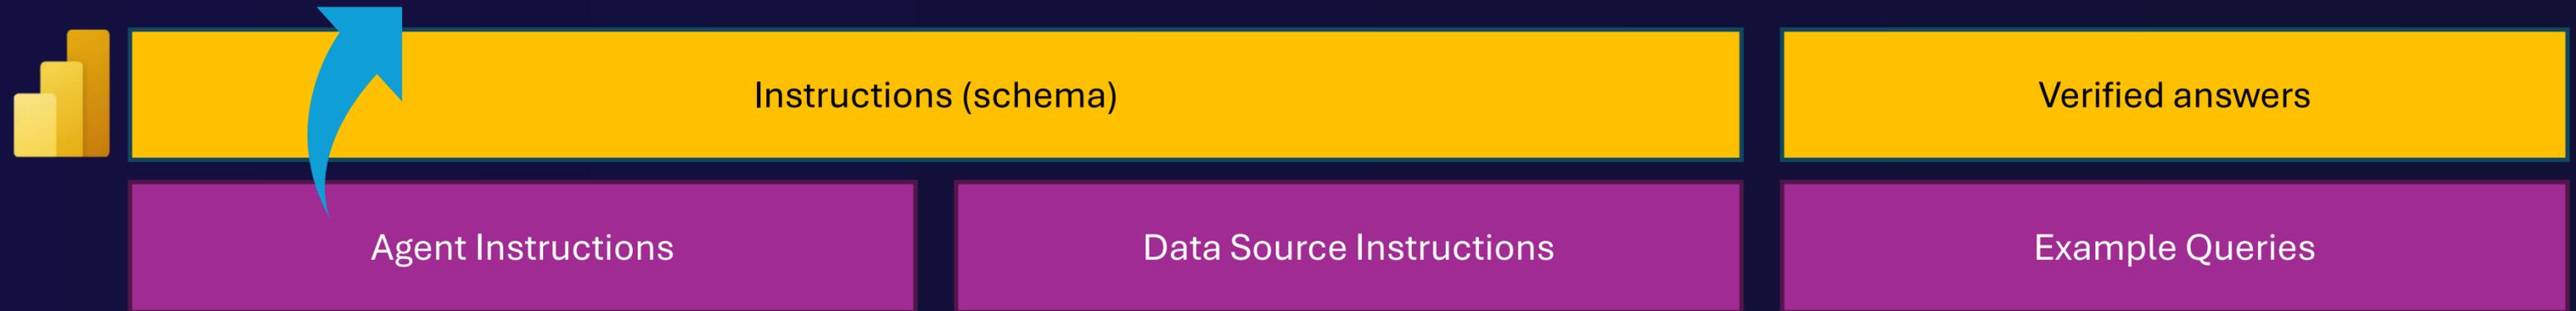


Data Agent building blocks

- Easy configurable
- Exposable
- Multi-source
- Developer-friendly



Creator configurations



Intended Use Cases:

- Specify which data sources to route for specific questions
- Guide how the agent interprets user intent

What to include:

- Map terminology to user intent
- Define steps for handling key questions
- Set response style and tone

Intended Use Cases:

- Provide clarity on how to use each source
- Guide table selection for topics
- Ensure correct query logic

What to include:

- Descriptions of key tables & columns
- Recommended table usage for certain topics
- Example column values and filter patterns

Intended Use Cases:

- Reflect advanced logic
- Examples should highlight join patterns or expected values

What to include:

- Representative set of questions a consumer may ask

Different types of supported sources



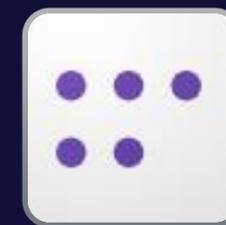
Lakehouse



Warehouse



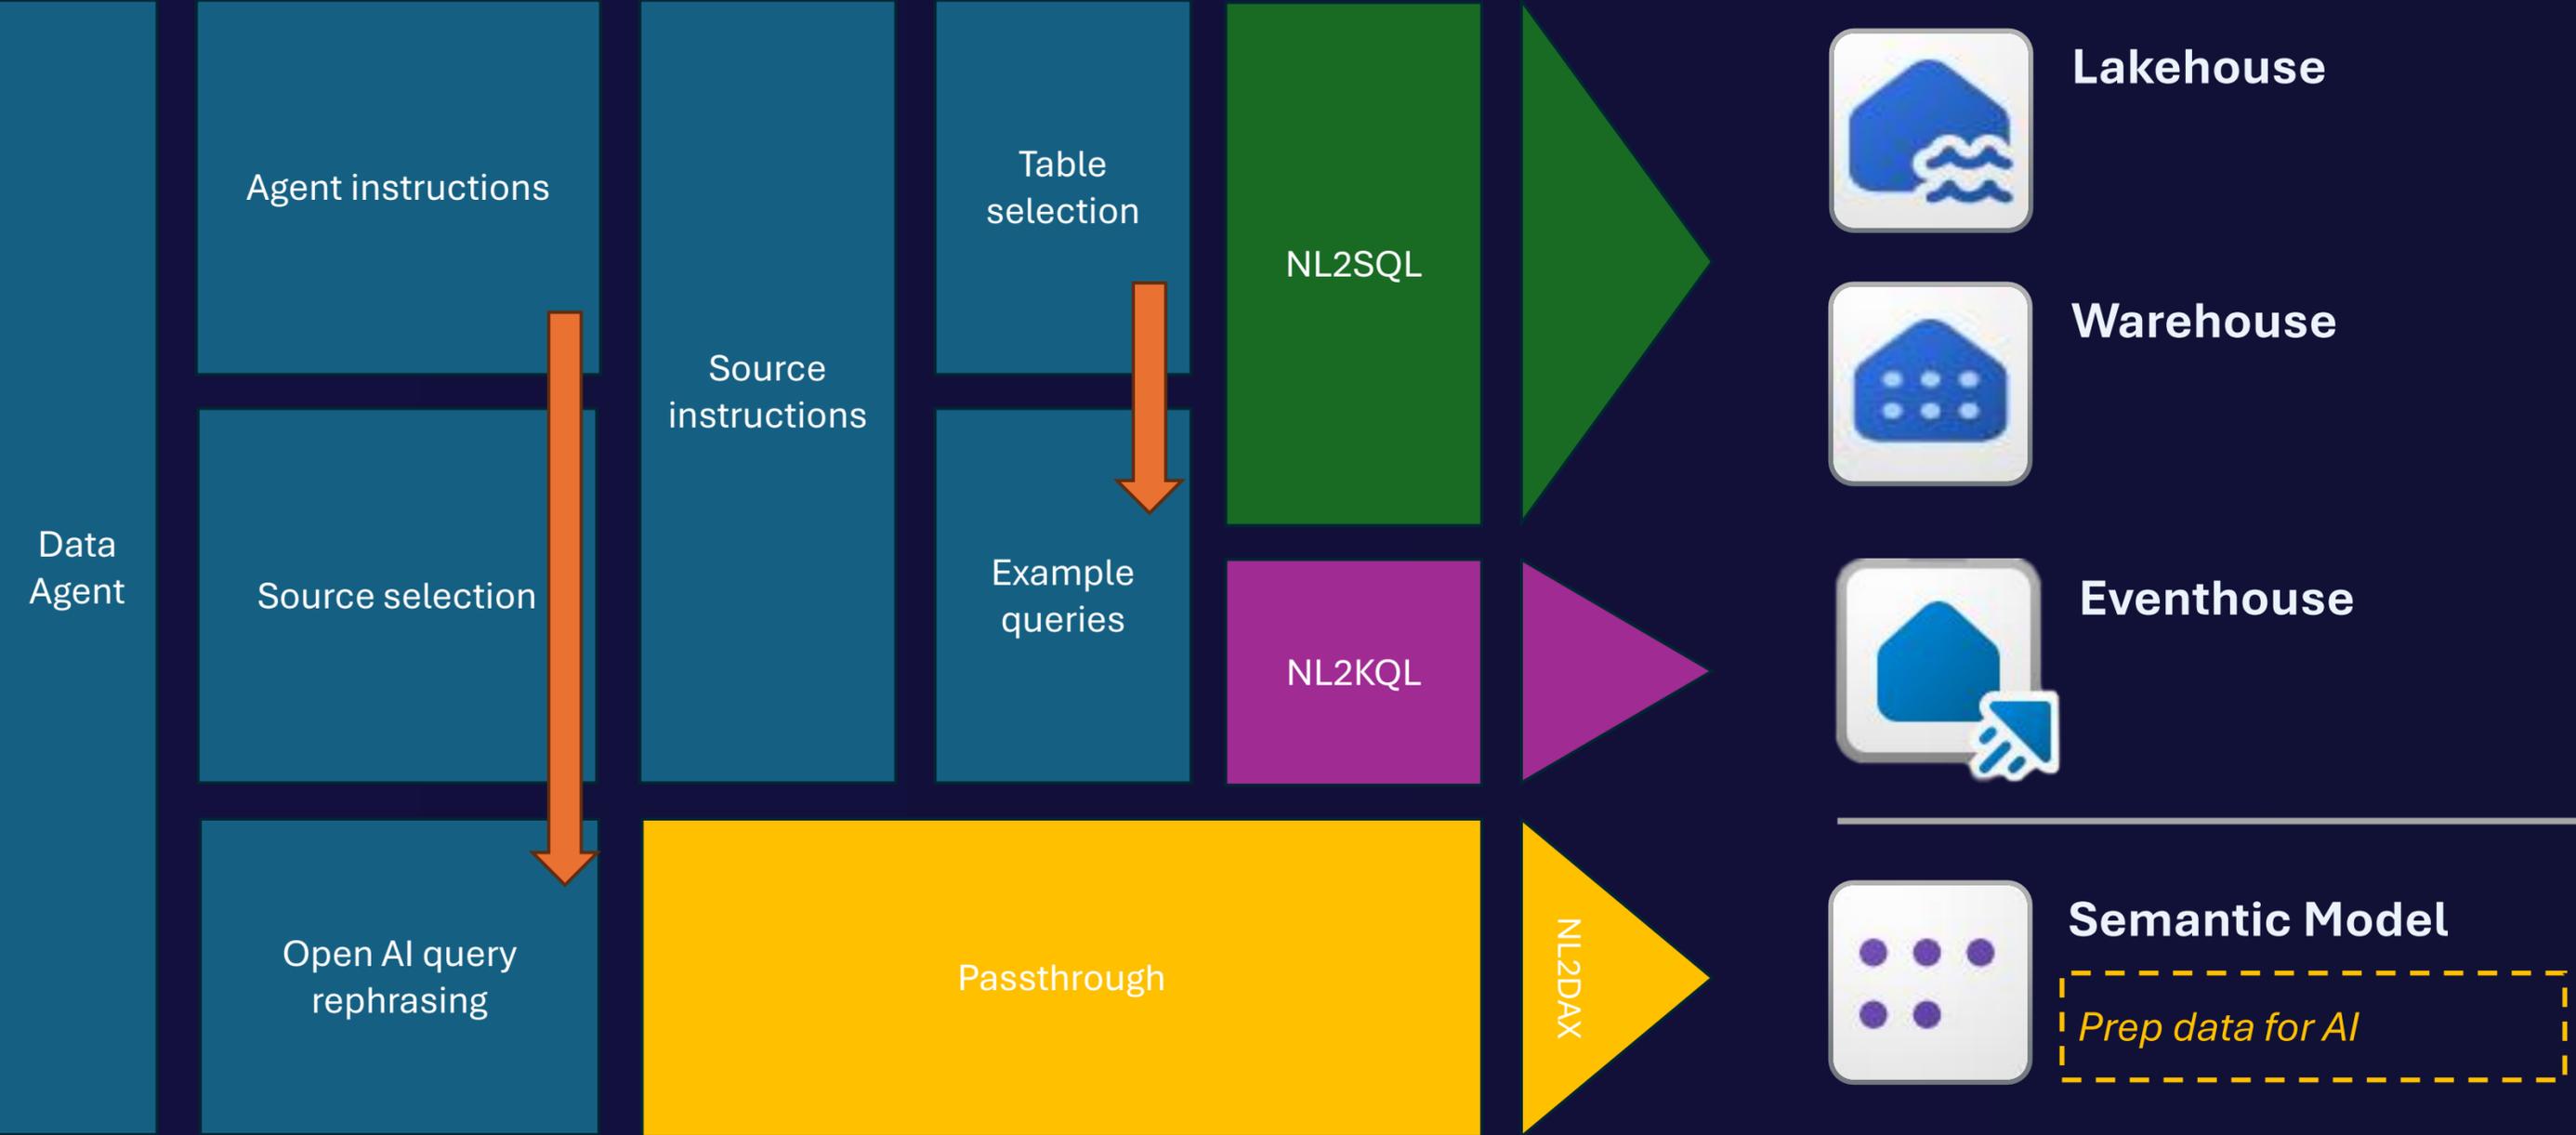
Eventhouse



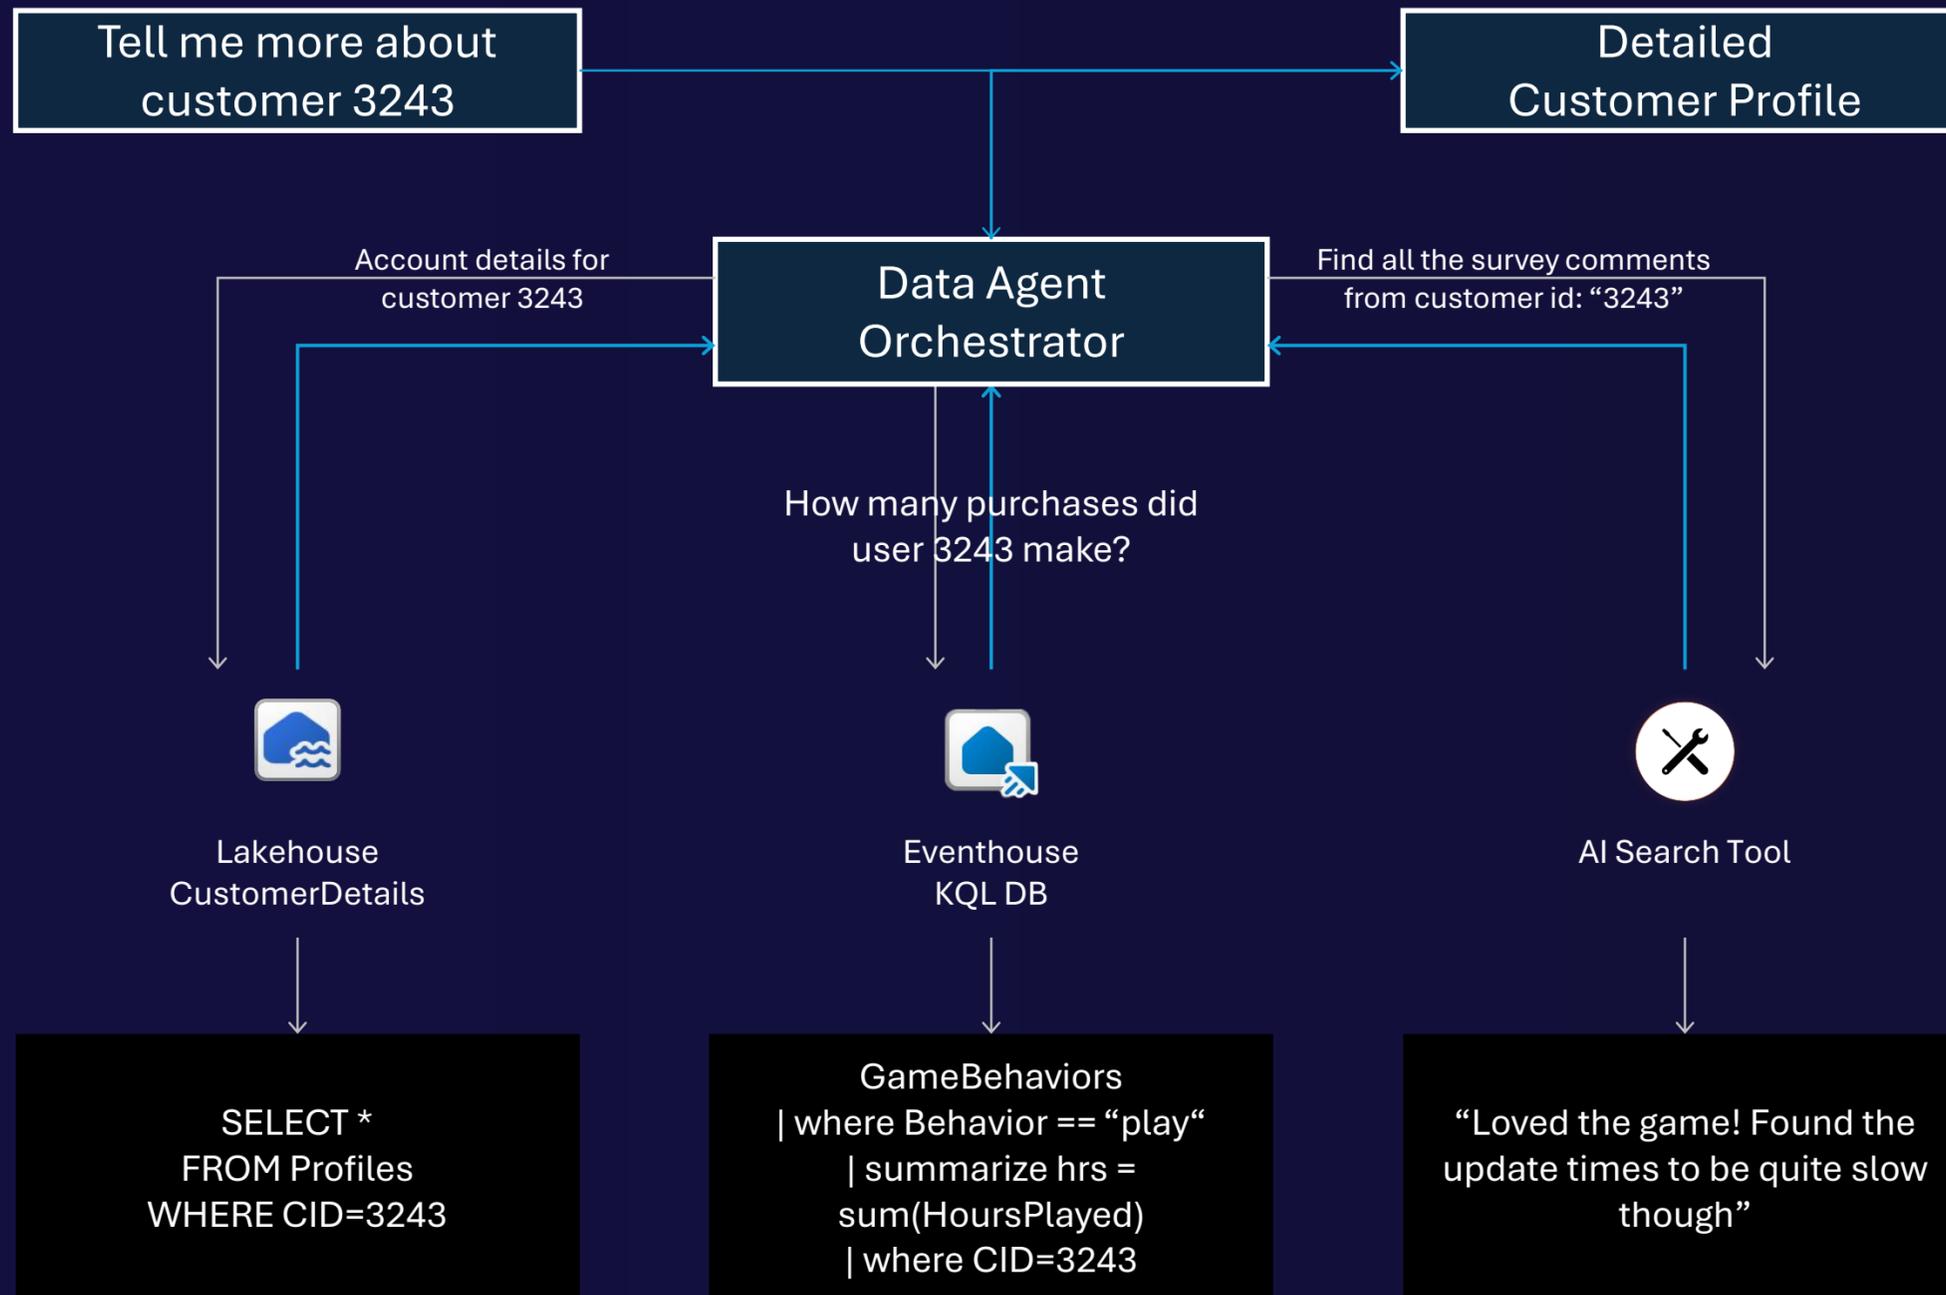
Semantic Model

And related sources (like mirrored sources)

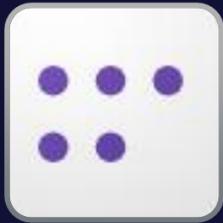
Data Agent orchestrator



The Data Agent flow



- 1 The data agent delegates questions to specific tools or data sources.
- 2 Data Sources are specialized at taking a natural language question and generating a response – query or result.
- 3 Data Sources & Tools supply their response back to the orchestrator – allowing it to process the information and think about the next step or return the summarized result.



Semantic models: Prep data for AI

Three main pillars

- Simplify schema
- Verified answers (links to visuals)
- AI Instructions

Prep data for AI

Get started

- Simplify the data schema
- Verified answers
- Add AI instructions

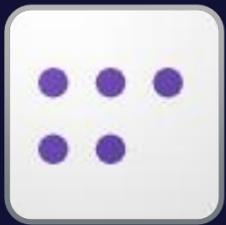
Prep this model to be AI-ready

Improve Copilot insights about this semantic model in a few simple steps. Publishing to the Power BI service will apply your changes to all Copilot chats, but updates may take a moment. [Learn more](#)

Simplify the data schema
Improve response accuracy by deselecting fields Copilot doesn't need to analyze.

Verified answers
Save predefined responses for your most critical business topics.

Add AI instructions
Help Copilot understand industry terms, business priorities, and important data fields.



Semantic models: Prep data for AI – Simplify schema

- Only keep what is relevant!
- Hide report-specific measures
 - Conditional formatting
 - Measures outputting texts like visual titles
- Disconnected tables likely cause issues
→ build relationships!

Keep in mind that Visual Calculations are not part of the semantic model!

Prep data for AI

Get started

Simplify the data schema

Verified answers

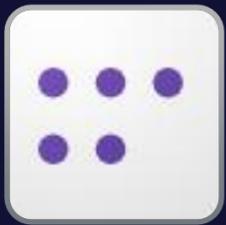
Add AI instructions

Simplify the data schema (preview)

Improve response accuracy by deselecting fields Copilot doesn't need to analyze. This creates a d

Filter by keyword

- ✓ Contoso Sales - end result
 - > _Measures
 - > Currency Exchange Rate
 - > Customer
 - > Date
 - ✓ Order
 - CurrencyCode
 - CustomerKey
 - DeliveryDate
 - Order Year
 - OrderDate
 - OrderKey
 - StoreKey
 - > Order Detail
 - > Product
 - > Sales
 - > Store



Semantic models: Prep data for AI – Verified answers

- Only keep what is relevant!
- Hide report-specific measures
 - Conditional formatting
 - Measures outputting texts like visual titles
- Disconnected tables likely cause issues
→ build relationships!

Keep in mind that Visual Calculations are not part of the semantic model!

Prep data for AI

Get started

Simplify the data schema

Verified answers

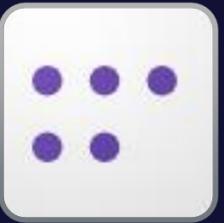
Add AI instructions

Simplify the data schema (preview)

Improve response accuracy by deselecting fields Copilot doesn't need to analyze. This creates a d

Filter by keyword

- Contoso Sales - end result
 - _Measures
 - Currency Exchange Rate
 - Customer
 - Date
 - Order
 - CurrencyCode
 - CustomerKey
 - DeliveryDate
 - Order Year
 - OrderDate
 - OrderKey
 - StoreKey
 - Order Detail
 - Product
 - Sales
 - Store



Semantic models: Prep data for AI – Verified answers

Prep data for AI

- Get started
- Simplify the data schema
- Verified answers**
- Add AI instructions

Verified answers (preview)

← All verified answers

Phrases connected to verified answers

Add phrases people might use when asking about the model data, and Copilot will respond with this visual. Test your phrases out in the chat before publishing.

 Show the top 3 categories by sales.

Copilot suggestions

Copilot uses AI. Always review content for mistakes. [Read terms](#)

Visual

The visual below will appear in Copilot's responses to phrases you add. Any changes to the report won't impact this visual.

CategoryName	Sales Amount
Computers	0.55bn
Cell phones	0.2bn
Home Appliances	0.15bn



Semantic models: Prep data for AI – AI instructions

- Unstructured guidance (textual)
- No guarantee the guidance is followed exactly – see it as a hint!
- Response formatting
 - Style
 - Tone
- Business specific logic – for example:
 - Define what is a top performer – high or low?
 - Peak / off-season: which months are these?

Prep data for AI

- Get started
- Simplify the data schema
- Verified answers
- Add AI instructions**

Add AI instructions (preview)

Provide context on important data fields, business goals, and industry terms to improve Copilot in

Add instructions for Copilot, such as:

- * For revenue-related questions, use the Sales Performance table.
- * EBITDA: Earnings Before Interest, Taxes, Depreciation, and Amortization.
- * If asked whether something is "on target", show Total Sales for the current month.

0/10000 characters

Keep in mind!

The answers will only be as good as your data!

siso

Shit In, Shit Out: If you're putting shit data into an experiment, you're going to get shit data out.

Your results are bubkes, what did you expect? SIS0!



Define success



Reliable answers



Secured



Well adopted

SDK

- **Programmatic Management:** Create, update, and delete Data Agent artifacts seamlessly.
- **Data Source Integration:** Easily connect to and integrate multiple data sources, for enhanced data analysis and insight generation.
- **OpenAI Assistants API Support:** Use the OpenAI Assistants API for rapid prototyping and experimentation.
- **Workflow Automation:** Automate routine tasks to reduce manual efforts and improve operational efficiency.
- **Resource Optimization:** Optimize configuration and management of Data Agent resources to better align with specific customer needs.

 Copy

```
%pip install fabric-data-agent-sdk
```

Diagnostics

“AI is wrong” to structured root cause analysis!

Where to start your investigation when the answer is not what you’re looking for?

Permissions

- RLS / CLS
- Item level access
- governance policies

Prompt & grounding

- Clear question
- Proper metadata

Source grounding

- Semantic model, lakehouse, warehouse, eventhouse
- Relationships and item naming

Notes for model

- Tone & voice
- Context (what is a good performer, high or low?)
- Hints (example queries)

Generated query analysis

- Diagnostics

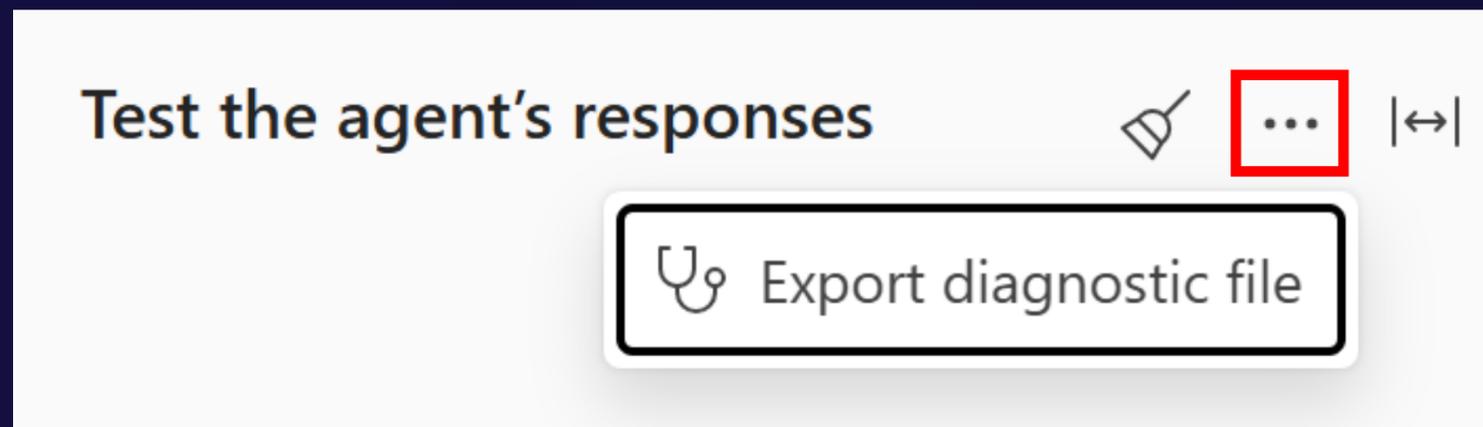
Capacity & performance

- Storage mode (for semantic models)
- Model size
- Load on the capacity

Diagnostics

“AI is wrong” to structured root cause analysis!

Where to start your investigation when the answer is not what you’re looking for?



User prompts

Data source schemas
and configurations

File metadata and
contents

Environment and
artifact identifiers
(GUIDS)

Any proprietary
information you may
have entered

Diagnostics

Source instructions

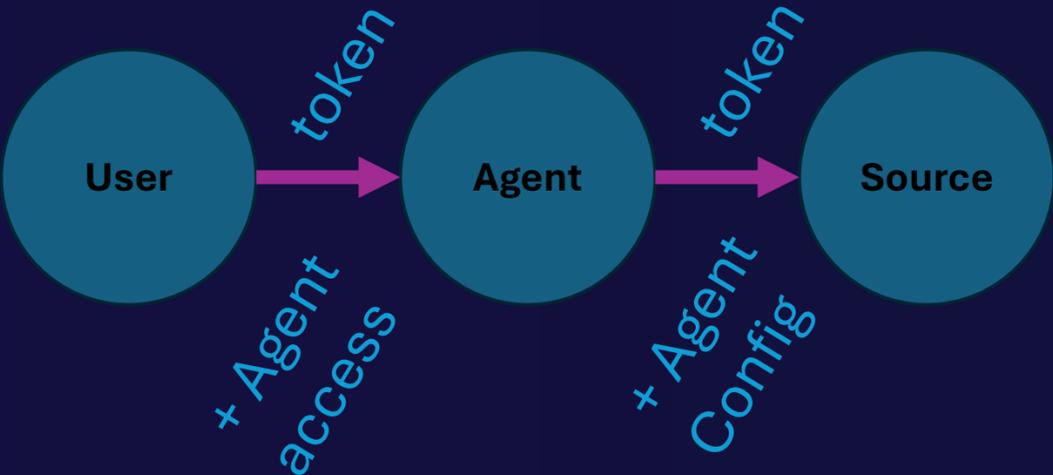
Schema information

...and much more!

```
{
  "downloaded_at": "2026-03-04T14:01:40.795Z",
  "rolloutEnvironment": "PROD",
  "stage": "sandbox",
  "artifactId": "a513749d-39bc-42d0-a002-dcd1a68d6811",
  "workspaceId": "073a733c-faf5-4e8e-b7f0-563480204837",
  "config": {
    "configuration": {
      "dataSources": [
        {
          "workspaceId": "073a733c-faf5-4e8e-b7f0-563480204837",
          "artifactId": "c4436ed1-8386-47bd-b142-b571c9675698",
          "type": "lakehouse_tables",
          "id": "c4436ed1-8386-47bd-b142-b571c9675698",
          "isSelected": true
        }
      ],
      "additionalInstructions": "The table FactInternetSales has a relationship to DimProduct based on the column ProductKey. \n\nWhen I ask for the sales amount, or total sales, or revenue",
      "schemaVersion": "https://developer.microsoft.com/json-schemas/fabric/item/dataAgent/definition/dataAgent/2.1.0/schema.json"
    },
    "etag": "\"0x8DE2B456941B0C8\""
  },
  "datasources": {
    "c4436ed1-8386-47bd-b142-b571c9675698": {
      "schema": {
        "dataSourceInfo": {
          "type": "lakehouse_tables",
          "lakehouse_id": "c4436ed1-8386-47bd-b142-b571c9675698",
          "lakehouse_name": "DemoData",
          "lakehouse_workspace_id": "073a733c-faf5-4e8e-b7f0-563480204837",
          "sql_endpoint": "xlhywod5jbtbjmffiaqc4qg5e-hrztub7v7khe5n7qky2iaicig4.datawarehouse.fabric.microsoft.com",
          "id": "c4436ed1-8386-47bd-b142-b571c9675698",
          "display_name": "DemoData",
          "user_description": "dbo.DimProduct.Color represents the product color. ",
          "additional_instructions": null,
          "metadata": {},
          "elements": [
            {
              "id": "dbdb800b-08f7-4157-9e1e-03c8b46d7ea7",
              "is_selected": false,
              "display_name": "dbo",
              "type": "lakehouse_tables.schema",
              "description": null,
              "children": [
                {
                  "id": "babfe0c3-a481-438a-a92b-d1bce5d19ab6",
                  "is_selected": true,
                  "display_name": "dbo_DimDate",
                  "type": "lakehouse_tables.table",
                  "description": null,
                  "children": [
                    {
                      "id": "552ae565-6aab-439b-9325-b2833b66a3c3",
                      "is_selected": true
                    }
                  ]
                }
              ]
            }
          ]
        }
      }
    }
  }
}
```

Permissions

Data access



Pass-through mechanics

- User credentials flow to the data source
- Permission evaluated per-user, per-query
- RLS/OLS etc. are baked-in

Service principals (SPN) are **NOT** supported, Access requires a user identity.

Data source permission requirements

Semantic model

Build permissions *(for now! Read permissions only soon)*

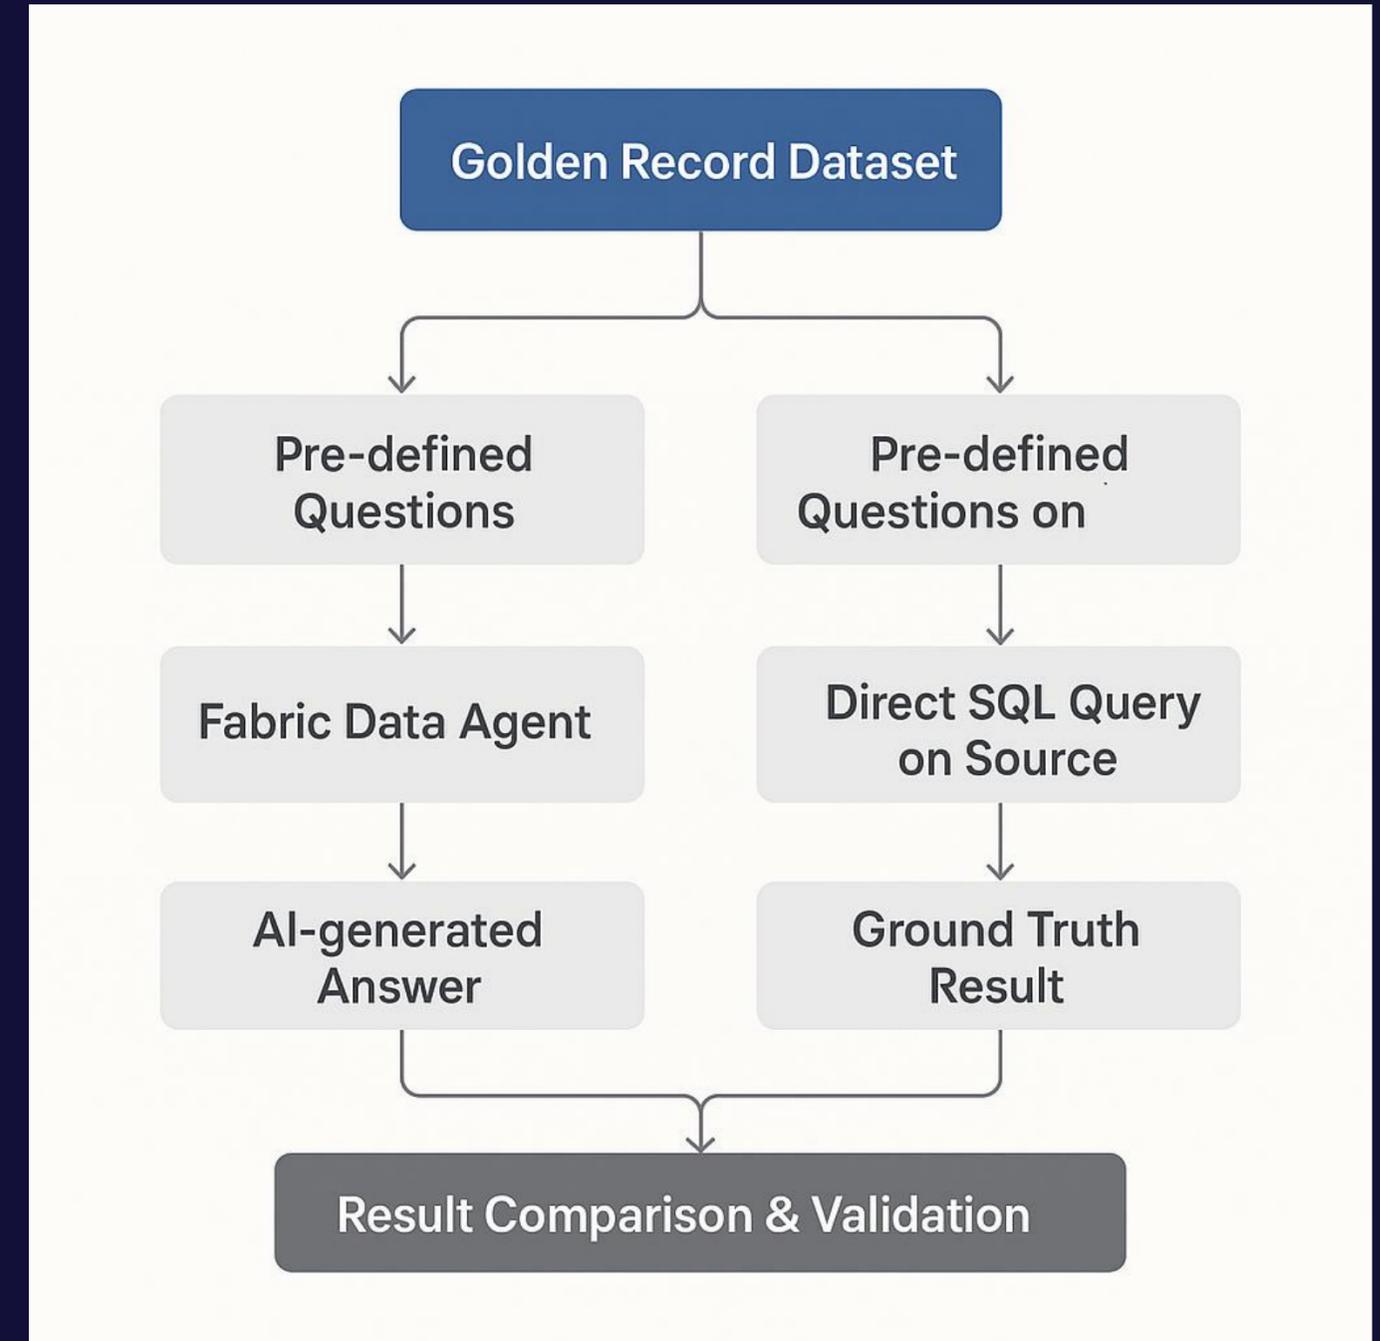
Agent access

Permission	Description	View Config	Edit Agent
Default (None)	N/A	✗	✗
View details	View the configuration and settings, but make no changes.	✓	✗
Edit and view details	Edit the configuration and settings.	✓	✓

Golden record testing

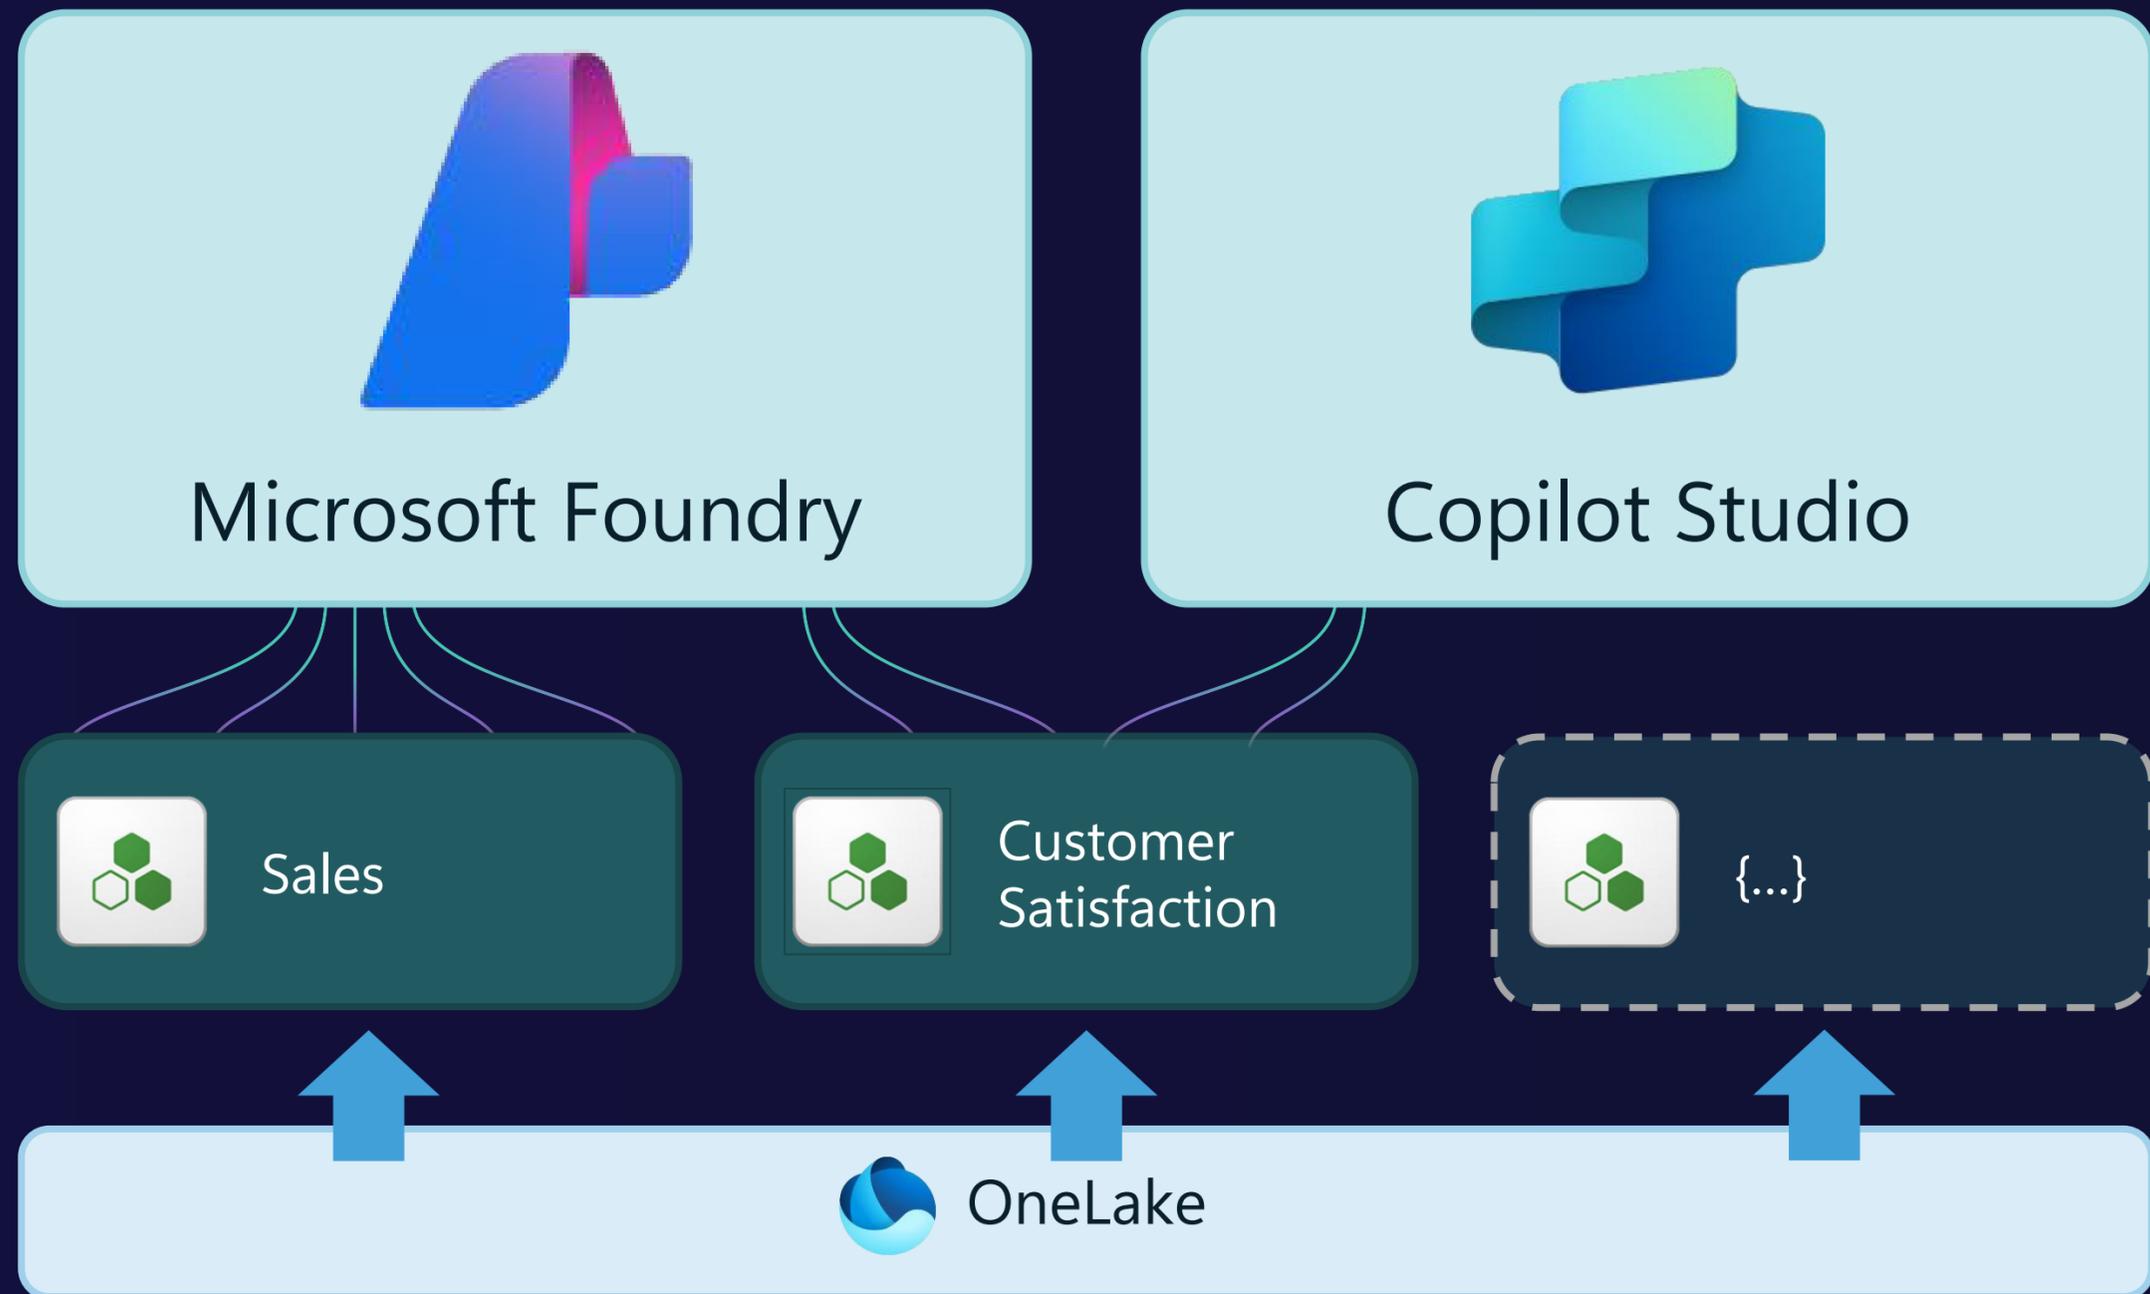
- Establishes trust by validating your Data Agent against a known, authoritative dataset.
- Results are compared one-to-one
- Enables objective evaluation

Gains confidence in the data agent, produces accurate, explainable and repeatable answers

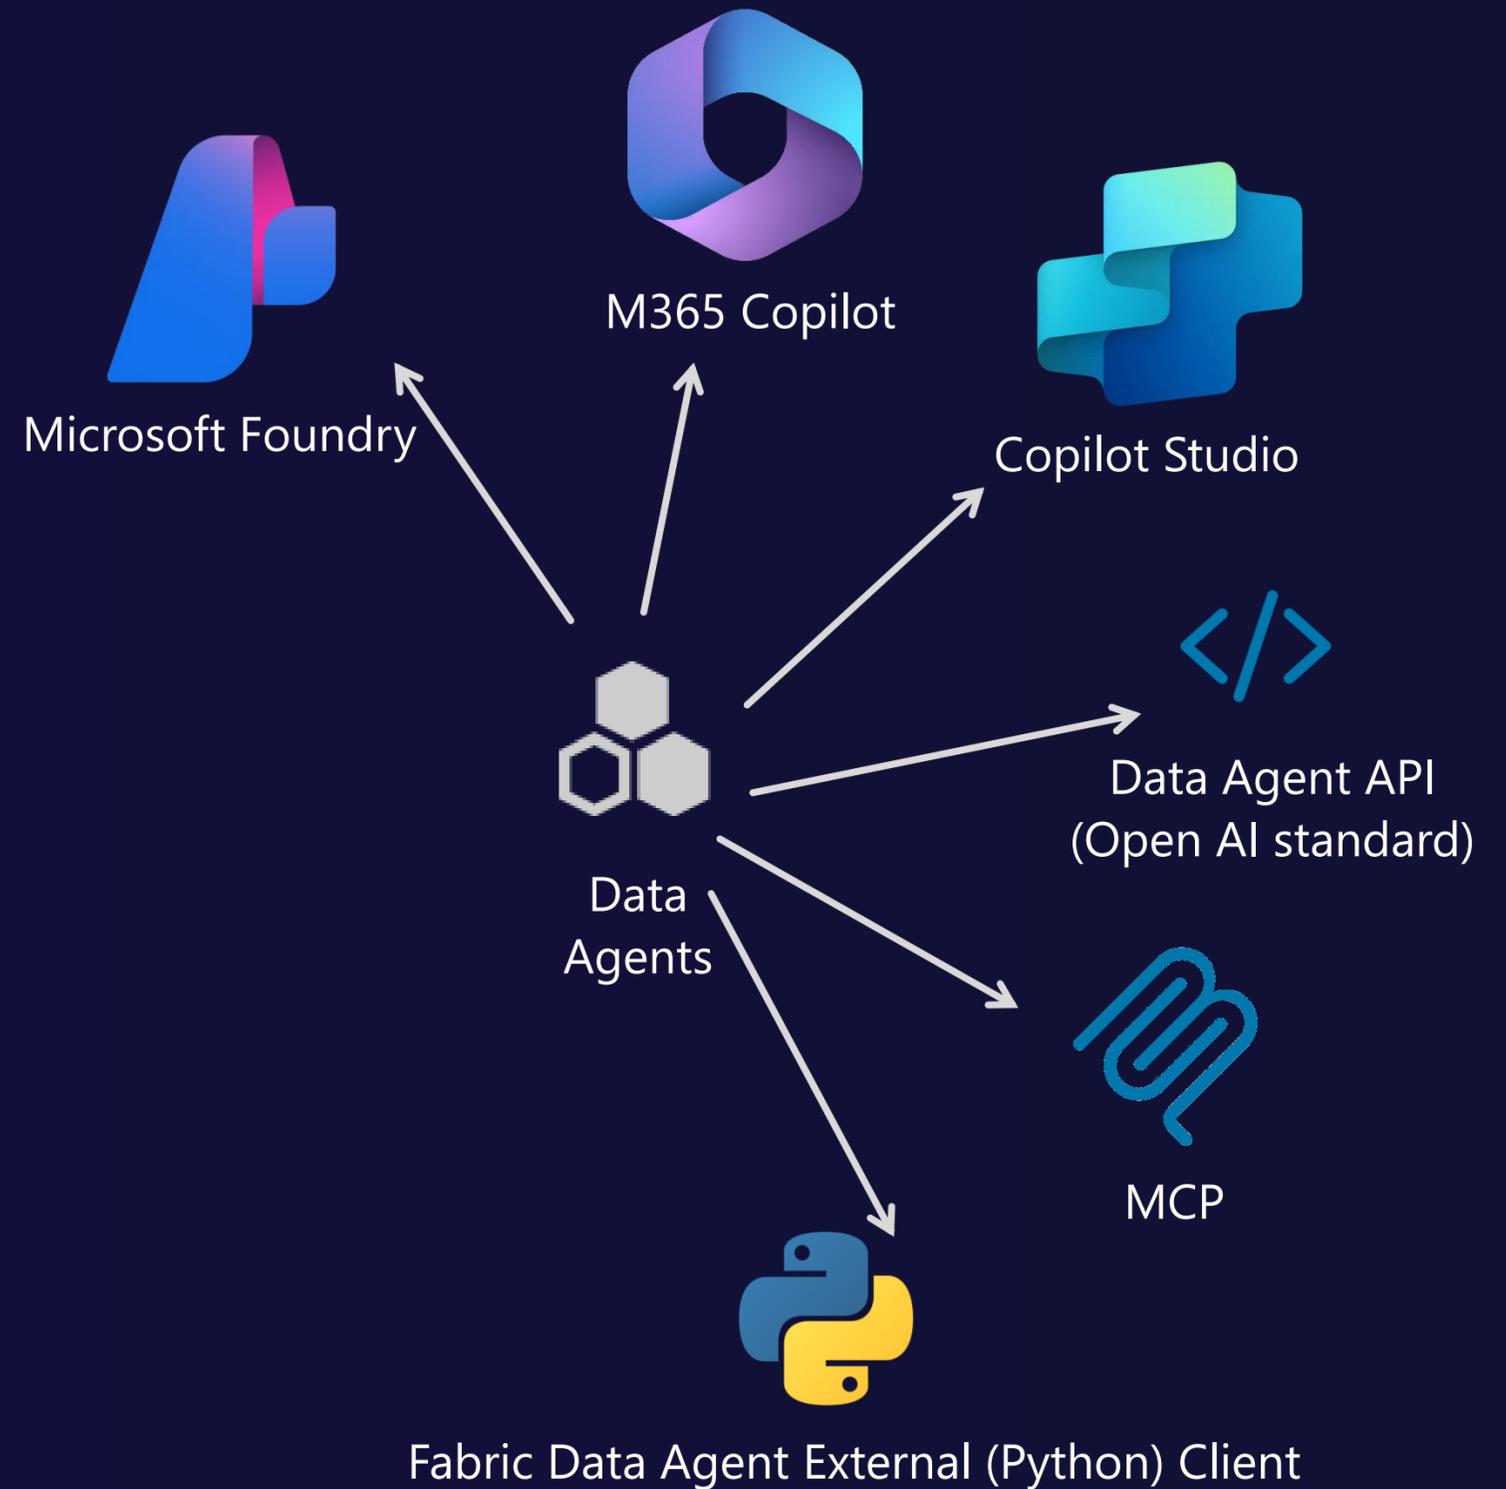


Demo

Agent Architecture One to rule them all?



Part of a larger ecosystem



Wrap-up

- Fabric Data Agents enable **natural, governed access** to data across OneLake, semantic models, lakehouses, warehouses, and eventhouses.
- **Strong foundations matter most**, with well-prepared schemas, verified answers, and clear AI instructions driving accuracy and trust.
- **Security and diagnostics are essential**, ensuring correct permissions, explainable behavior, and structured root-cause analysis when results differ.
- Data Agents fit into **a broader ecosystem**, integrating with Fabric, Power BI, Copilot Studio, Microsoft Foundry, and custom applications.

Key takeaway:

Reliable AI answers come from reliable data, thoughtful design, and continuous validation.

How was the session?



Complete Session Surveys in
Whova for your chance to WIN
PRIZES!

